



#7
sp
12-31-03

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

In re the Application of:

TI-26019

Iain Robertson

Art Unit: 2182

Serial No: 09/615,312

Examiner: Eron J. Sorrell

Filed: July 13, 2000

Conf. No.: 6993

For: Using Write Request Queue to Prevent Bottlenecking on Slow Ports in Transfer
Controller With Hub and Ports Architecture

TRANSMITTAL LETTER ACCOMPANYING CERTIFIED COPY OF
PRIORITY APPLICATION UNDER 35 U.S.C. §119

RECEIVED

DEC 24 2003

Technology Center 2100

Commissioner for Patents
P.O. Box 1450
Alexandria, VA 22313-1450

MAILING CERTIFICATE UNDER 37 C.F.R. §1.8(a)
I hereby certify that the above correspondence is being deposited
with the U.S. Postal Service as First Class Mail in an envelope
addressed to: Commissioner for Patents, P.O. Box 1450,
Alexandria, VA 22313-1450 on December 18, 2003.


Robin E. Bamum

Dear Sir:

Submitted herewith is a certified copy of **European** Patent Application No. **9916706.6**,
filed on **July 16, 1999**, in the **European** Patent Office and from which priority under 35 U.S.C.
§119 is claimed for the above-identified application.

Respectfully submitted,



Robert D. Marshall, Jr.
Attorney for Applicant
Reg. No. 28,527

Texas Instruments Incorporated
P.O. Box 655474, MS 3999
Dallas, TX 75265
(972) 917-5290

This Page Blank (uspto)



INVESTOR IN PEOPLE

11/12
21/6/15, 132
CIC-26019
The Patent Office
Concept House
Cardiff Road
Newport
South Wales
NP10 8QQ

CERTIFIED COPY OF PRIORITY DOCUMENT


I, the undersigned, being an officer duly authorised in accordance with Section 74(1) and (4) of the Deregulation & Contracting Out Act 1994, to sign and issue certificates on behalf of the Comptroller-General, hereby certify that annexed hereto is a true copy of the documents as originally filed in connection with the patent application identified therein.

I also certify that the attached copy of the request for grant of a Patent (Form 1/77) bears an amendment, effected by this office, following a request by the applicant and agreed to by the Comptroller-General.

In accordance with the Patents (Companies Re-registration) Rules 1982, if a company named in this certificate and any accompanying documents has re-registered under the Companies Act 1980 with the same name as that with which it was registered immediately before re-registration save for the substitution as, or inclusion as, the last part of the name of the words "public limited company" or their equivalents in Welsh, references to the name of the company in this certificate and any accompanying documents shall be treated as references to the name with which it is so re-registered.

In accordance with the rules, the words "public limited company" may be replaced by p.l.c., plc, P.L.C. or PLC.

Re-registration under the Companies Act does not constitute a new legal entity but merely subjects the company to certain additional company law rules.



Signed

JEvens.

Dated 06 September 2000

This Page Blank (uspto)

9916706.6

19 JUL 1999
19 JUL 99 E462586-7 002849
P41/7710 0.01 - 9916706.6

Request for grant of a patent

(See the notes on the back of this form. You can also get an explanatory leaflet from the Patent Office to help you fill in this form)

The Patent Office

Cardiff Road
Newport
South Wales
NP9 1RH

1. Your reference

TI-26019 GB

2. Patent application number

(The Patent Office will fill in this part)

9916706.6

3. Full name, address and postcode of the or of each applicant (underline all surnames)

TEXAS INSTRUMENTS LIMITED

800 Pavilion Drive
Northampton NN4 7YL
United Kingdom

Patents ADP number (if you know it)

If the applicant is a corporate body, give the country/state of its incorporation

ENGLAND

00766378004

4. Title of the invention

IMPROVEMENTS IN OR RELATING TO MICROPROCESSORS

5. Name of your agent (if you have one)

HOLT, Michael

"Address for service" in the United Kingdom to which all correspondence should be sent (including the postcode)

European Patent Activity
P.O. Box 5069
Northampton NN4 7ZE
United Kingdom



Patents ADP number (if you know it)

07663826004

6. If you are declaring priority from one or more earlier patent applications, give the country and the date of filing of the or of each of these earlier applications and (if you know it) the or each application number

Country

Priority application number
(if you know it)

Date of filing
(day / month / year)

7. If this application is divided or otherwise derived from an earlier UK application, give the number and the filing date of the earlier application

Number of earlier application

Date of filing
(day / month / year)

8. Is a statement of inventorship and of right to grant of a patent required in support of this request? (Answer 'Yes' if:

a) any applicant named in part 3 is not an inventor, or

b) there is an inventor who is not named as an applicant, or

c) any named applicant is a corporate body.

See note (d))

Yes

Patents Form 1/77

9. Enter the number of sheets for any of the following items you are filing with this form. Do not count copies of the same document

Continuation sheets of this form

Description	18			
Claim(s)	10	Sheets/	29 ?	Claims
Abstract	1			
Drawing(s)	9 + 9 + 9 0	Sheets/	10 ?	Figures

10. If you are also filing any of the following, state how many against each item.

Priority documents	0
Translations of priority documents	
Statement of inventorship and right to grant of a patent (Patents Form 7/77)	1
Request for preliminary examination and search (Patents Form 9/77)	1
Request for substantive examination (Patents Form 10/77)	0
Any other documents (please specify)	0

11. I/We request the grant of a patent on the basis of this application.

Signature	Date
Michael Cowles - Company Secretary	15/07/1999

12. Name and daytime telephone number of person to contact in the United Kingdom
Zakia Ryan - +44 (0) 1604 663361

Warning

After an application for a patent has been filed, the Comptroller of the Patent Office will consider whether publication or communication of the invention should be prohibited or restricted under Section 22 of the Patents Act 1977. You will be informed if it is necessary to prohibit or restrict your invention in this way. Furthermore, if you live in the United Kingdom, Section 23 of the Patents Act 1977 stops you from applying for a patent abroad without first getting written permission from the Patent Office unless an application has been filed at least 6 weeks beforehand in the United Kingdom for a patent for the same invention and either no direction prohibiting publication or communication has been given, or any such direction has been revoked.

Notes

- If you need help to fill in this form or you have any questions, please contact the Patent Office on 0645 500505.
- Write your answers in capital letters using black ink or you may type them.
- If there is not enough space for all the relevant details on any part of this form, please continue on a separate sheet of paper and write "see continuation sheet" in the relevant part(s). Any continuation sheet should be attached to this form.
- If you have answered 'Yes' Patents Form 7/77 will need to be filed.
- Once you have filled in the form you must remember to sign and date it.
- For details of the fee and ways to pay please contact the Patent Office.

IMPROVEMENTS IN OR RELATING TO MICROPROCESSORS

TECHNICAL FIELD OF THE INVENTION

The present invention relates generally to the field of microprocessors and more specifically to the field of digital signal processor devices. The techniques, described, however, are applicable even more generally to digital devices for performing data exchange between ports of a multiple port device.

BACKGROUND OF THE INVENTION

The present invention deals with the write request queue as applied to the transfer controller with hub and ports architecture.

The transfer controller with hub and ports is the subject of U.K. Patent Application number 9909196.9, filed April 21, 1999. The transfer controller with hub and ports is a significant basic improvement in data transfer techniques in complex digital systems and provides, along with many other features, for the uniform implementation of port interfaces at the periphery of such systems.

Some of these ports may be slow, (i.e. they may have relatively low throughput). Others may be relatively fast, having the throughput and speed of a current central processing unit.

Some bottlenecks in data transfer operations can result from this disparity of performance within the device or system ports. To clarify the problem solved by the teachings of the present application, it is helpful to review the characteristics, architecture, and functional building blocks of the transfer controller with hub and ports. While direct memory access (DMA) is a powerful tool in a digital signal processor system, they have their limitations. The fundamental limitation of a conventional direct memory access engine is that adding additional channel capacity requires additional hardware (in general, a replication of a complete channel). Some optimizations can be made in this area, such as sharing registers between multiple channels, but in general, the following rule holds: N-channels costs N times as much as a single channel.

This Page Blank (uspto)

Conventional direct memory access techniques initiate reads from a source, and subsequently pass the data on to a destination.

The source transfers will initially proceed at full rate. However, if the source has higher data transfer bandwidth than the destination, this data will backlog within the direct memory access engine. This will eventually slow the rate at which source transfers are issued. Thus the source data transfer bandwidth is effectively restricted to that of the destination. If another channel has a different source port and a different destination port, there are no conflicts using the conventional read driven approach. However, if the source port of the other channel is the same, the other channel could not be processed. This makes for inefficiency. In a device that supports only one transfer at a time, this is acceptable. However, the transfer controller with hub and ports device supports multiple concurrent transfers and other provisions must be made. A normal transfer process in the known art starts by reading data from the source and then writing it to the destination. The source read drives the process in that it occurs first, and everything follows as a consequence.

With a conventional read driven approach, the source will start reading data which will be passed to the destination. However, if the destination is slow, a backlog of data waiting to be written will eventually cause the source read process to stall because it will not have anywhere to put the data read. With only one channel this is acceptable, but if there are multiple channels, conflicts occur. The source for this channel is stalled and cannot respond to more read requests. However, it is desirable to be able to service a different channel instead.

These basic limitations to conventional data transfer techniques led to the initial development of the transfer controller with hub and ports. The transfer controller with hub and ports is a unique mechanism which consolidates the functions of a direct memory access and other data movement engines in a digital signal processor system (for example, cache controllers) into a single module.

Consolidation of such functions has both advantages and disadvantages. The most important advantage of consolidation is that it will, in general, save hardware since multiple instantiations of the same type of address generation hardware will not have to be implemented.

On a higher level, it is also advantageous to consolidate address generation since it inherently makes the design simpler to modify from a memory-map point of view. For example, if a peripheral is added or removed from the system, a consolidated module will be the only portion of the design requiring change.

In a distributed address system (multi-channel direct memory access for example), all instances of the direct memory access channels would change, as would the digital signal processor memory controllers.

Fundamental disadvantages of the consolidated model, however, are its inherent bottlenecking, resulting from conflicting multiple requests, and its challenge to higher clock rates. Additionally, there is in general an added complexity associated with moving to a consolidated address model, just because the single module is larger than any of the individual parts it replaces.

The transfer controller with hub and ports, to which the present application relates, is a highly parallel and highly pipelined memory transaction processor. This transfer controller with hub and ports serves as a backplane to which many peripheral and/or memory ports may be attached.

SUMMARY OF THE INVENTION

The write request queue disclosed herein provides a means to prevent bottlenecking which would otherwise occur in a device capable of processing multiple data transfer requests simultaneously. This bottlenecking results from the occurrence of blocking transfers, (i.e. when a transfer between ports A and B prevents a transfer from ports A and C from occurring). The teachings of the present application rely on write driven processing and provides a write request queue to store transfer

requests which must be temporarily held while the transfer in process completes.

The transfer controller with hub and ports also includes a mechanism to eliminate this feature on ports which are fast enough to handle writes at up to the full speed of the transfer controller with hub and ports. In general, only the internal memory port of the transfer controller with hub and ports can support such writes at up to the full speed of the transfer controller with hub and ports. The transfer controller with hub and ports is independent of such writes, as it does not distinguish between the internal and the external ports and what physically resides on them. It simply performs read or write driven processing via a tie off at each port interface.

In write driven processing, the destination write drives the process. First the hub engine requests the destination to reserve some space for the data to be written. When this request is granted, the source read can occur. The actual destination write will occur in due course. The ability to store these requests in the write request queue supplies this capability.

BRIEF DESCRIPTION OF THE DRAWINGS

The present invention will now be further described, by way of example, with reference to certain exemplary embodiments illustrated in the accompanying drawings in which:

Figure 1 illustrates a block diagram of the basic principal features of the transfer controller with hub and ports;

Figure 2 illustrates the functional blocks of the transfer controller hub and its interface to external ports and internal memory port master;

Figure 3 illustrates two transfers operating concurrently, one between ports A and B (higher priority), and the second between ports A and C (lower priority);

Figure 4 illustrates the flow of the write driven process in which the data does not necessarily arrive in the same order as the requests;

Figure 5 illustrates two write requests, the first in time proceeding from a port with higher latency A and the second in time proceeding from a port of lower latency B and the need to resolve the protocol when the data from the second request port B arrives first;

Figure 6 illustrates the three part structure of the external ports interface units;

Figure 7 illustrates a functional block diagram showing conceptually the hub interface unit and the write request queue;

Figure 8 illustrates a single integrated circuit including multiple processors and the transfer controller of this invention;

Figure 9 illustrates a block diagram form an example of one of the multiple processors illustrated in Figure 8; and

Figure 10 illustrates further details of the very long instruction word digital signal processor core illustrated in Figure 9.

DETAILED DESCRIPTION OF PREFERRED EMBODIMENTS

Figure 1 illustrates a block diagram of the basic features of the transfer controller with hub and ports. The transfer controller with hub and ports is basically a data transfer controller which has at its front end portion, a queue manager 100 for receiving, prioritizing, and dispatching data in the form of transfer request packets. The queue manager 100 connects within the hub unit 110 to the channel registers 120. Channel registers 120 receive the data transfer request packets and process them first by prioritizing them and assigning them to one of the N channels. Each channel represents a priority level. These channel registers 120 interface with the source control pipeline 130 and destination control pipeline 140. These are address calculation units for source (read) and destination (write) operations.

Outputs from these pipelines are broadcast to M Ports (six shown in Figure 1 as 150 through 155). The ports 150 to 155 are clocked either at the main processor clock frequency or at a lower external device clock frequency. Read data from one port, (e.g.

port 150), having a destination write address of port 153 is returned to the hub destination control pipeline through the routing unit.

The transfer controller with hub and ports disclosed herein, introduces several new ideas supplanting the previous transfer controller technology. First, it is uniformly pipelined. In the previous transfer controller designs, the pipeline was heavily coupled to the external memory type supported by the device. In the preferred embodiment, the transfer controller with hub and ports contains multiple external ports, all of which look identical to the hub. Thus peripherals and memory may be freely interchanged without affecting the transfer controller with hub and ports. Secondly, the transfer controller with hub and ports concurrently executes transfers. That is, up to N transfers may occur in parallel on the multiple ports of the device, where N is the number of channels in the transfer controller with hub and ports core.

Each channel in the transfer controller with hub and ports core is functionally just a set of registers. These registers track the current source and destination addresses, the word counts and other parameters for the transfer. Each channel is identical, and thus the number of channels supported by the transfer controller with hub and ports is highly scalable. Thirdly, the transfer controller with hub and ports includes a mechanism for queuing transfers up in a dedicated queue RAM.

The transfer controller with hub and ports and its closely associated functional units, may be broken into five main entities (refer to Figure 2):

- (1) The request bus master 200 input takes in transfer request packets from the transfer request feed mechanism 201. These transfer requests originate from processor elements or other devices. These transfer request packets are input to the queue manager 220 which is within the transfer controller with hub and ports. A transfer request is a command (in the preferred embodiment the command word is of quad word length or double long word length, i.e. approximately 128 bits in length) to the transfer

controller with hub and ports to move a specified number of data elements from one global address to another.

(2) The transfer controller with hub and ports, shown within the dashed line of Figure 2, includes queue manager 220, dual port RAM 225 queue storage, and the hub engine 235 with channel registers 239, source/destination pipelines and routing unit 249.

Also illustrated are the channel parameters registers 234 and port parameters registers 236. The dashed line dissects the external ports interface units 240-247, indicating that the transfer controller with hub and ports proper, terminates within these ports. The hub engine 235 performs the servicing of the transfers, breaking them into smaller transfers that the devices can handle.

(3) The internal memory port (IMP) node master 250 is a special transfer controller with hub and ports port, interfaces to the data transfer bus (DTB) 255. The internal memory port node master 250 connects a port representing internal memory to the processor nodes. Processor nodes interfaces here include all distributed internal data memory and internal program memory controllers and to all the other control registers on the processor system.

(4) The data transfer bus (DTB) 255 connects the port representing internal memory to the memory interfaces in the digital signal processor nodes and other units in the core of the chip.

(5) The external ports interfaces 240 through 247 act as a buffer between the hub engine and the external memory or peripheral port.

Consider an example where a high priority transfer occurs between fast main external memory and a slow peripheral, such as a serial port and a lower priority transfer between the main external memory and the internal memory also occurs. Write driven processing is used to prevent occurrences of a backlog at a slow serial port which would result in a main memory read stall. In this situation according to the prior art the main memory would be able to pass data only at the same effective bandwidth as the slow

serial port. During this stall, a lower priority external to internal memory transfer would not be able to proceed until the slow serial port transfer was complete.

With write driven processing, reads for data to be transferred to the serial port will be scheduled only if the serial port has room for them. The external memory will not need to stall because it has fetched data and has no location to which the data can be sent. As soon as the serial port can absorb no more data, the hub will be free to schedule transfers for the lower priority external to internal memory transfer until space for another transfer becomes available in the serial port. As the data transfer bandwidth of the serial port may be several orders of magnitude less than the external memory, a significant improvement in efficiency occurs.

Consider the following prior art example illustrated in Figure 3. There are two transfers operating concurrently, one between ports A 300 and ports B 301, and another between ports A 300 and ports C 302. The transfer between port A 300 and port B 301 has a higher priority 304 than the transfer between port A 300 and port C 302. The interface at port A 300 has a higher bandwidth 305 than both the bandwidth 306 of port B 310 and the bandwidth 307 of port C 302. When the transfer from port A 300 to port B 301 starts up, data read from port A 300 backlogs within the device because port B 301 is slower than port A 300. New requests to port A 300 can only be submitted at the rate that port B 301 is clearing the backlog. The interface at port A 300 cannot be used for transfers between port A 300 and port C 302 during the times that the interface at port A 300 cannot be used for fetching data for port B 301. This is a blocking transfer. When transferring data from port A 300 to port B 301, the prior art approach is to read from port A 300, and when data is returned, pass it on to port B 301.

However, if the interface at port A 300 has a higher bandwidth than that of port B 301, this can result in a large backlog of data requested from port A 300 which cannot yet be passed on to port B 301. If the transfer controller is multitasking, this backlog can

prevent the interface at port A 300 from being applied to another transfer, such as the lower priority from port A 300 to port C 302 concurrently as illustrated in Figure 3.

The use of a write request queue in conjunction with write driven processing prevents this blocking problem. According to the present teachings, a write to port B 301 is requested before reading from port A 300. Data is read from port A 300 at the rate required to fulfill the data transfer rate of port B 301. No backlog occurs, and multitasking can occur effectively.

The flow chart of Figure 4 illustrates how the write request queue mechanism disclosed herein prevents such bottlenecks. In step 401, the transfer controller hub receives a data transfer request to move data from port A to port B. The transfer controller hub then interrogates port B to determine if it can accept a transfer of the required data size (step 402). If it can, the destination is told to expect the data and the space required for the data is marked as in use. This is in anticipation that the data won't actually arrive until later clock cycles (step 403).

Once the transfer controller hub has established that the destination has reserved space for the data, the source read proceeds (step 404). When the data arrives, it is passed directly to the destination, port B (step 405). In step 406, port B forwards the data and frees up the space ready for the next transfer (step 407).

On the other hand, if the destination does not have space available, the device can attempt a different transfer instead.

In step 412 a check is made to see if a transfer from port A to port C is pending. If so, and if port C can accept a transfer, then steps 403 through 407 are duplicated as steps 413 through 417 relative to port C. Note that this process should be repeated for all other possible destinations with port A as the source. If any transfers are pending from port A, then these may be serviced during any interval when port B cannot receive data.

The transfer controller hub includes a mechanism to limit the rate at which source data is read to the rate at which it can

subsequently be written without introducing congestion in the data transfer engine (i.e. the hub). A write reservation station is needed on all ports, except those as fast as the transfer controller hub. In the preferred embodiment of the transfer controller with hub and ports, the fastest port is the internal memory port but any port could potentially be the fastest.

Generally a write reservation station is required for any port which has a write request queue. Whenever an entry is placed in the write request queue, the data for that write will arrive sometime later. On arrival this data goes into the reservation station. It is called a reservation station rather than a write data queue because the data does not necessarily arrive in the same order as the requests. This will be illustrated below.

Consider the example illustrated in Figure 5. A first channel 501 is set up for a transfer between port A 502 and port C 504.

A second channel 507 is set up for a transfer between port B 503 and port C 504. Port B 503 has a lower read latency 508 than the read latency of port A 502. That is, the time between a read request being made and data being returned is lower for port B 503 than port A 502. The first channel 501 services a write request to port C 504. Space is reserved for write data, and the request is granted. A second channel 507 services a write request from port B 503 to port C 504. Again, space is reserved for write data and request is granted. The second channel 507 write request occurred after first channel 501 write request, the data from the second channel 507 read from port B 503 may arrive at port C 504 before data from first channel 501 read from port A 502 because of lower latency of port B 503. Thus the arrival order cannot be used to associate the data with the requests made earlier.

Write driven processing is applied only at the hub interface unit (HIU). This occurs at the generic or common front end of the external port interface units. This is where the reservation station and the write request queue will reside. Here, the request for the write will be made before the data to be written is present. This data will arrive later. However, this can all be

hidden from application unit (AU), which is the back end of the external ports interface unit. The hub interface unit will not make a request to write to the application unit until the data arrives, at which point a conventional write request with data available can be made to the application unit.

Figure 6 illustrates the three divisions of the external port interface units into generic front end (hub unit interface 601), synchronizer 602 and customized back end (application unit interface 603). The write request queue and write reservation station is contained wholly within the front end portion hub interface 601.

Figure 7 illustrates the write request queue as a portion of the generic front end hub interface unit. The write request queue within the dot-dashed lines consists of address and status registers 703 for write reservation station '0' 713 and address and status registers 704 for write reservation station '1' 714. The datapath of the write request queue consists of potentially a greater number of write reservation entries. In the example of Figure 7 only two write reservation station entries are provided. The first is labeled write reservation station '0' 713 and the second is labeled write reservation station '1' 714.

The sequence in a write driven processing transfer of data proceeds as follows. Note that the hub interface unit portion of the port illustrated in Figure 7 is the destination port in this example. The data transfer bus pipeline issues a pre-write command to the destination port. This is step 402 in Figure 4, where the transfer controller hub interrogates the destination port to determine if it can accept a transfer of the required size. In Figure 7, this proceeds to the address and status registers only, and is labeled a primary write. In executing this command, the hub interface unit stores an indication of the transfer, such as the channel number, data size, source or destination address and sequence number, in the address and status register 703 that is associated with write reservation station 713. Thus address and status register 703 serves as an identifier register for the

corresponding write reservation station 713. If the transfer controller hub receives a 'yes' response indicating that the port can receive the data, this destination port is told to expect the data. The space required for the data in the write reservation station is marked as in use in anticipation of this data although the data won't actually arrive until later clock cycles.

Once the transfer controller hub has established that the destination has reserved space in write reservation station 713 for the data, the source read proceeds. This occurs at another port, the source port. The source pipeline output stage 700 issues a read command to the source port. When the data arrives via the path through the response queue (a path identical to response queue 701 in the source port) it is passed through the routing unit 710 and a secondary write is issued by the destination pipeline output stage 710. This secondary write contains the data to be written.

This data is first written to write reservation station 713 in the reserved space in the destination port and then is passed from the hub interface unit as write data. Once this secondary write data is passed to the write request queue and write reservation station of the destination port, it is forwarded to the external application unit. Then space in the write reservation station is freed up and ready for the next transfer (steps 405 and 406 of Figure 4). As noted earlier in Figure 4, if the destination did not have space available, the transfer controller with hub and ports can attempt a different transfer instead.

When the hub interface unit accepts a pre-write command, it stores identifier data, such as channel number, data size, source or destination address or sequence number for example, in the address and status register 703 associated with write reservation station 713. When the transfer controller hub destination pipeline output 710 sends the write data to the port, the port checks the identifier data and determines the location to which to write the corresponding write reservation station entry. In this example, the associated write reservation station is write reservation station 713.

In the preferred embodiment the write request queue has a number of other requirements. The address and 2-bit size information are stored in a first in first out (FIFO) queue. A counter within the write reservation station counts the data received and compares it to the pre-write indicated data size. When the data received is equal to the pre-write data size, the data transfer from the transfer controller hub to the hub interface unit is completed and the data is ready to be transferred to the external device, also referred to as the application unit.

In summary, the data transfer is as follows. The hub interface unit performs a read operation and receives the data from an internal memory port (if the source is an internal memory device) or from another hub interface unit (if the source is a peripheral port application unit) via the transfer controller hub routing unit. It then performs a write function, transferring data to another of the transfer controller peripheral ports via the destination hub interface unit.

With write driven processing, data reads from a source port will be scheduled only if the destination port has room for the data. The source port will not need to stall because it has fetched data and has no location to send the data. As soon as the destination port can absorb no more data, the transfer controller hub will be free to schedule transfers for the lower priority transfers until space for another transfer becomes available in the original destination port. In this example the original destination port is a serial port. As the data transfer bandwidth of the serial port may be several orders of magnitude less than the data transfer bandwidth of other ports such as internal or external memory, a significant improvement in efficiency occurs. The method and apparatus disclosed herein use a write request queue in conjunction with write driven processing. This prevents this problem by requesting the write to the destination port before reading from the source port. It results in data being read from the source at the rate required to fulfill the destination port.

No backlog occurs, and multitasking can occur effectively.

Figure 8 illustrates from a higher level an overview of an multiprocessor integrated circuit employing the transfer controller with hub and ports. There are four main functional blocks. The transfer controller with hub and ports 110 and the ports including ports external port interface units 240 to 247 and internal memory port 250 are the first two main functional blocks. The other two main functional blocks are the transfer request feed mechanism 201 and the data transfer bus (DTB) 255. These are closely associated functional units that are but not a part of the transfer controller with hub and ports 110. Transfer request feed mechanism 201 is coupled to plural internal memory port nodes 870, 871 and 872. Each of these internal memory port nodes includes an independently programmable data processor, which may be a digital signal processor, and corresponding cache memory or other local memory.

The internal construction of these internal memory port nodes is not important. For the purpose of the method and apparatus disclosed herein it sufficient that each of the internal memory port nodes can submit transfer requests via transfer request feed mechanism 201 and has memory that can be a source or destination for data. Transfer request feed mechanism 201 prioritizes these packet transfer requests in a manner not relevant to this invention. Transfers originating from or destined for internal memory port nodes 870, 871 or 872 are coupled to transfer controller with hub and ports 110 via data transfer bus 255 and internal memory port master 250. As previously described, internal memory port master 250 may not require the write driven process of this invention if internal memory port nodes 870, 871 and 872 have memory transfer bandwidth equivalent to the memory transfer bandwidth of transfer controller with hub and ports 110. Figure 8 highlights the possible connection of data transfer bus 255 to multiple internal memory port nodes 870, 871 and 872 and the possible connection of multiple transfer request nodes to transfer request feed mechanism 201. This represents an example of the mode of use of the write driven process of this invention and not its only context of use.

Figure 9 illustrates a block diagram of an example of a preferred processor and cache memory combination implementing the internal memory nodes 870, 871 and 872 of Figure 8. Each internal memory node 870, 871 and 872 preferably includes a digital signal processor core and corresponding instruction and data cache memory.

Transfer controller with hub and ports 110 provides for all data communication among internal memory nodes 870, 871 and 872, external input/output (I/O) devices and peripherals at external ports 240 to 247, and internal memory at local memory port 252.

Each internal memory node 870, 871 and 872 preferably comprises a very long instruction word (VLIW) digital signal processor core 44, program memory controller (PMC) 46, data memory controller (DMC) 48, an emulation, test, analysis and debug block 50, local memory and data transfer bus (DTB) interface 52. Internal memory nodes 870, 871 and 872 and transfer controller with hub and ports 110 communicate over a pair of high throughput buses. Transfer request feed mechanism 201 is used by digital signal processor cores 44 to specify and request transactions in transfer controller with hub and ports 110. Data transfer bus (DTB) 255 is used to load and store data from objects in the global memory map. While any given digital signal processor core 44 can access its own internal local memory within the cluster without permission from transfer controller with hub and ports 110, any access to global memory outside of its local memory requires a transfer controller directed data transfer, whether the access is to external memory or to another digital signal processor local memory. The overall architecture is scalable, allowing for the implementation of many internal memory nodes, although three is currently the preferred embodiment. It should be noted that architectural details, such as the number of digital signal processor cores, and their instruction set architectures are not essential to the invention. This microprocessor architecture is exemplary only, and the present teachings are applicable to many microprocessor architectures.

Figure 10 is a block diagram illustrating more detail of digital signal processor core 44 illustrated in Figure 9. Digital signal processor core 44 is a 32-bit eight-way VLIW pipelined processor. The instruction set consists of fixed length 32-bit reduced instruction set computer (RISC) type instructions that are tuned for digital signal processing applications. Almost all instructions perform register-to-register operations and all memory accesses are performed using explicit load/store instructions. As shown in Figure 10, instruction pipeline 58 consists of fetch stage 60 and decode stage 62. Fetch stage 60 retrieves program codes into the processor core from instruction cache 64 under control of program memory controller 46 in groups of eight instructions called a fetch packet. Decode stage 62 parses the fetch packet, determines parallelism and resource availability and constructs an execute packet of up to eight instructions. Each instruction in the execute packet is then translated into control signals to drive the appropriate units in execution pipeline 66. Execution pipeline 66 consists of two symmetrical datapaths, datapath A 68 and datapath B 70, a common 64-bit load/store unit group D-unit group 72, and a common branch unit group P-unit group 74. Each datapath contains 32-word register file (RF) 76, and four execution unit groups, A-unit group 78, C-unit group 80, S-unit group 82, and M-unit group 84. Overall there are ten separate unit groups in execution pipeline 66. Eight of these units may be scheduled concurrently every cycle. Each functional unit group contains plural functional units, some of which are duplicated between unit groups. In total there are nine 32-bit adders, four 32-bit shifters, three boolean operators, and two 32 bit by 16 bit multipliers. The multipliers are each configurable into two 16 bit by 16 bit multipliers or into four 8 bit by 8 bit multipliers. The memory at internal memory nodes 870, 871 and 872 is preferably partitioned between instruction cache memory 64 controlled via program memory controller 46 and data cache memory and random access memory 88 controlled via data memory controller 48. These

memory partitions are employed by digital signal processor core 44 in a conventional manner.

Each digital signal processor core 44 may request data transfers in several ways. Digital signal processor core 44 may issue a data transfer request to transfer controller with hub and ports 110 in response to an explicit data transfer instruction.

The data transfer instruction must specify the data source, the data destination and the data amount. These specifications may be by immediate fields in the instructions or by parameters stored in registers or memory. It is preferable that each digital signal processor core 44 be capable of requesting any data transfer that can be serviced by transfer controller with hub and ports 110. Thus any digital signal processor core 44 may transfer data internally or externally and load or read any internal memory node.

Each digital processor core 44 preferably also includes automatic mechanisms for generating requests for data transfer for cache service. Thus an instruction cache miss preferably causes program memory controller 46 to generate a data transfer request from another data source to fill a line of instruction cache 64 with data including program instructions stored at the address generating the cache miss. Similarly, a data cache miss on a data read preferably causes data memory controller 48 to generate a data transfer request to retrieve data to fill a line in data cache/random access memory 88 with corresponding data. These instruction and data are stored in a higher level of memory. This higher level of memory may be an on-chip combined cache used by all digital signal processor cores 44 or it may be external to the multiprocessor integrated circuit. There are two alternatives for data cache misses on data writes. In a write through mode, a data write by digital processor core 44 that misses data cache/random access memory 88 causes data memory controller 48 to generate a data transfer request to store the write data in the appropriate location in a higher level of memory. In a writeback mode, a data write by digital processor core 44 that misses data cache/random access memory 88 causes data memory controller 48 to generate a

data transfer request to recall corresponding data in the appropriate location from a higher level of memory for storage in data cache/random access memory 88. The write data is then written into data cache/random access memory 88 overwriting the corresponding data just recalled from the higher level of memory.

This process is referred to as write allocation within the data cache.

Data memory controller 48 preferably also employs a data transfer request to handle data writeback to a higher level memory upon cache eviction of a dirty entry. A dirty cache entry includes data that has been modified since it was recalled from a higher level of memory. This modified data corresponds to a later state of the program than the data stored in the higher level of memory.

When such data must be replaced to make room for new cache data, referred to as cache eviction, this dirty data must be written back to the higher level of memory to maintain the proper program state.

Transfer controller with hub and ports 110 is preferably employed for this writeback of evicted dirty cache entries.

WHAT IS CLAIMED IS:

1. A method of data transfer between a source port and a destination port of a transfer controller with plural ports, said method comprising the steps of:

in response to a data transfer request, querying said destination port to determine if said destination port is capable of receiving data of a predetermined size;

if said destination port is not capable of receiving data, waiting until said destination port is capable of receiving data,

if said destination port is capable of receiving data, reading data of said predetermined size from said source port and transferring said read data to said destination port.

2. The method of claim 1 further comprising:

providing at least one write reservation station for each port.

3. The method of Claim 1 or Claim 2, wherein said step of querying said destination port includes:

determining whether any write reservation station of said destination port has not been allocated for receipt of data, and

if at least one write reservation is not allocated for receipt of data, determining said destination port can receive data and allocating a write reservation station for receipt of data.

4. The method of claim 3, further comprising:

transferring data from a write reservation station storing data to be transferred to an application unit coupled to said destination port at a data transfer rate of said application unit; and

deallocating said write reservation station upon transfer of data to said application unit.

5. The method of claim 3, wherein:

said step of allocating a write reservation station includes storing a data identifier corresponding to said write reservation station; and

said step of transferring said read data to said destination port includes storing said read data in a write reservation station having a data identifier corresponding to said read data.

6. The method of any preceding claims, further comprising:

while waiting until said destination port is capable of receiving data, determining if a second data transfer is pending between said source port and a second destination port, and

if a second data transfer is pending, servicing said second data transfer.

7. The method of claim 6, wherein:

said step of servicing said second data transfer includes querying said second destination port to determine if said second destination port is capable of receiving data of said predetermined size;

if said second destination port is not capable of receiving data, waiting until said second destination port is capable of receiving data,

if said second destination port is capable of receiving data, reading data of said predetermined size from said source port and transferring said read data to said second destination port.

8. A method of data transfer between a source port and a destination port of a transfer controller with plural ports substantially, as described herein with reference to any of the accompanying drawings.

9. A data transfer controller comprising:

a request queue controller for receiving, prioritizing and dispatching data transfer requests, each data transfer request specifying a data source, a data destination and a data quantity to be transferred;

a data transfer hub connected to request queue controller for effecting dispatched data transfer requests;

a plurality of ports each having an interior interface connected to said data transfer hub and an exterior interface configured for an external memory/device expected to be connected to said port, said interior interface and said exterior interface operably connected for data transfer therebetween; and

said data transfer hub arranged for controlling data transfer from a source port corresponding to said data source to a destination port corresponding to said data destination in a quantity corresponding to said data quantity to be transferred of a currently executing data transfer request, said data transfer hub further controlling said source port and said destination port to

in response to a data transfer request, query said destination port to determine if said destination port is capable of receiving data of a predetermined size;

if said destination port is not capable of receiving data, waiting until said destination port is capable of receiving data, and

if said destination port is capable of receiving data, reading data of said predetermined size from said source port and transferring said read data to said destination port.

10. The data transfer controller of claim 9, wherein:

each port includes at least one write reservation station for storing data prior to transfer to said corresponding external memory/device;

said data transfer hub further controlling said destination port to determine whether any write reservation station of said destination port has not been allocated for receipt of data, and

if at least one write reservation is not allocated for receipt of data, determining whether said destination port can receive data and allocating a write reservation station for receipt of data.

11. The data transfer controller of claim 10, wherein:

said data transfer hub is further arranged for controlling said destination port to transfer data from a write reservation station to said corresponding external memory/device at a data transfer rate of said external memory/device, and

disallocating said write reservation station upon transfer of data from said write reservation station to said external memory/device.

12. The data transfer controller of claim 10 or claim 11, wherein:

each of said plurality of hubs further includes an identifier register corresponding to each write reservation station; and

said data transfer hub is further arranged for controlling said destination port to allocate a write reservation station by writing identifier data in said corresponding identifier register, and to store said read data in a write reservation station having a corresponding identifier stored in said identifier register corresponding to said write reservation station.

13. The data transfer controller of any of claims 9 to 12, wherein:

said data transfer controller further capable of servicing a second transfer request between said source port and a second

destination port while waiting until said destination port is capable of receiving data.

14. The data transfer request controller of claim 13, wherein:
said data transfer controller is further arranged for controlling said second destination port to query said second destination port to determine if said second destination port is capable of receiving data of said predetermined size;

if said second destination port is not capable of receiving data, waiting until said second destination port is capable of receiving data,

if said second destination port is capable of receiving data, reading data of said predetermined size from said source port and transferring said read data to said second destination port.

15. A data transfer controller substantially as described herein with reference to any of the accompanying drawings.

16. A data processing system comprising:

a plurality of data processors, each data processor capable of generating a data transfer request;

a request queue controller connected to said plurality of data processors, said request queue controller arranged for receiving, prioritizing and dispatching data transfer requests, each data transfer request specifying a data source, a data destination and a data quantity to be transferred;

a data transfer hub connected to request queue controller for effecting dispatched data transfer requests;

a plurality of ports, each of said plurality of ports having an interior interface connected to said data transfer hub identically configured for each port and an exterior interface configured for an external memory/device expected to be connected to said port, said

interior interface and said exterior interface operably connected for data transfer therebetween; and

said data transfer hub controlling data transfer from a source port corresponding to said data source to a destination port corresponding to said data destination in a quantity corresponding to said data quantity to be transferred of a currently executing data transfer request, said data transfer hub further arranged for controlling said source port and said destination port to

in response to a data transfer request, query said destination port to determine if said destination port is capable of receiving data of a predetermined size;

if said destination port is not capable of receiving data, waiting until said destination port is capable of receiving data, and

if said destination port is capable of receiving data, reading data of said predetermined size from said source port and transferring said read data to said destination port.

17. The data processing system of claim 16, wherein:

each port includes at least one write reservation station for storing data prior to transfer to said corresponding external memory/device;

said data transfer hub further arranged for controlling said destination port to determine whether any write reservation station of said destination port has not been allocated for receipt of data, and

if at least one write reservation is not allocated for receipt of data, determining said destination port can receive data and allocating a write reservation station for receipt of data.

18. The data processing system of claim 17, wherein:

said data transfer hub is further arranged for controlling said destination port to transfer data from a write reservation station to said corresponding external memory/device at a data transfer rate of said external memory/device, and

disallocate said write reservation station upon transfer of data from said write reservation station to said external memory/device.

19. The data processing system of claim 17, wherein:

each of said plurality of hubs further includes an identifier register corresponding to each write reservation station; and

said data transfer hub is further arranged for controlling said destination port to allocate a write reservation station by writing identifier data in said corresponding identifier register, and

store said read data in a write reservation station having a corresponding identifier stored in said identifier register corresponding to said write reservation station.

20. The data processing system of any of claims 16 to 19, wherein:

said data transfer controller is further capable of servicing a second transfer request between said source port and a second destination port while waiting until said destination port is capable of receiving data.

21. The data processing system of claim 20, wherein:

said data transfer controller is further arranged for controlling said second destination port to

query said second destination port to determine if said second destination port is capable of receiving data of said predetermined size;

if said second destination port is not capable of receiving data, waiting until said second destination port is capable of receiving data,

if said second destination port is capable of receiving data, reading data of said predetermined size from said source port and transferring said read data to said second destination port.

22. The data processing system of any of claims 16 to 21, further comprising:

a system memory connected to a predetermined one of said plurality of ports; and

wherein each of said data processors includes an instruction cache for temporarily storing program instructions for controlling said data processor, said data processor arranged for generating a data transfer for program cache fill from said system memory upon a read access miss to said instruction cache.

23. The data processing system of any of claims 16 to 22, further comprising:

a system memory connected to a predetermined one of said plurality of ports; and

wherein each of said data processors includes a data cache for temporarily storing data employed by said data processor, said data processor arranged for generating a data transfer for data cache fill from said system memory upon a read access miss to said data cache.

24. The data processing system of any of claims 16 to 23, further comprising:

a system memory connected to a predetermined one of said plurality of ports; and

wherein each of said data processors includes a data cache for temporarily storing data employed by said data processor, said data

processor arranged for generating a data transfer for data writeback to said system memory upon a write miss to said data cache.

25. The data processing system of any of claims 16 to 24, further comprising:

a system memory connected to a predetermined one of said plurality of ports; and

wherein each of said data processors includes a data cache for temporarily storing data employed by said data processor, said data processor arranged for generating a data transfer for write data allocation from said system memory to said data cache upon a write miss to said data cache.

26. The data processing system of any of claims 16 to 25, further comprising:

a system memory connected to a predetermined one of said plurality of ports; and

wherein each of said data processors includes a data cache for temporarily storing data employed by said data processor, said data processor arranged for generating a data transfer for data writeback to said system memory upon eviction of dirty data from said data cache.

27. The data processing system of any of claims 16 to 26, wherein:

said plurality of data processors, said request queue controller, said data transfer hub and said plurality of ports are disposed on a single integrated circuit.

28. The data processing system of any of claims 16 to 27, further comprising:

a data memory having a data transfer bandwidth on the same order as a data transfer bandwidth of said data transfer hub;

a second port connected to said data transfer hub and said data memory; and

said data transfer hub further arranged for controlling said source port and said destination port to not query said second port to determine if said destination port is capable of receiving data of a predetermined size if said second port is a destination port of a transfer request.

29. A data processing system substantially as herein described with reference to any of the accompanying drawings.

IMPROVEMENTS IN OR RELATING TO MICROPROSESSORS

ABSTRACT

A data transfer technique between a source port and a destination port of a transfer controller with plural ports. In response to a data transfer request (401), the transfer controller queries the destination port to determine if it can receive data of a predetermined size (402). If the destination port is not capable of receiving data, the transfer controller waits until said destination port is capable of receiving data (412). If the destination port is capable of receiving data, the destination port allocates a write reservation station to the data (403). Then the transfer controller reads data of the predetermined size from the source port (404) and transfers this read data to the destination port (405). The destination port forwards this data to an attached application unit, which may be memory or a peripheral, and then disallocates the write reservation station freeing space for further data transfer (406). This write driven process permits the transfer controller hub to service other data transfers from a fast source without being blocked by a slow destination.

Figure 4.

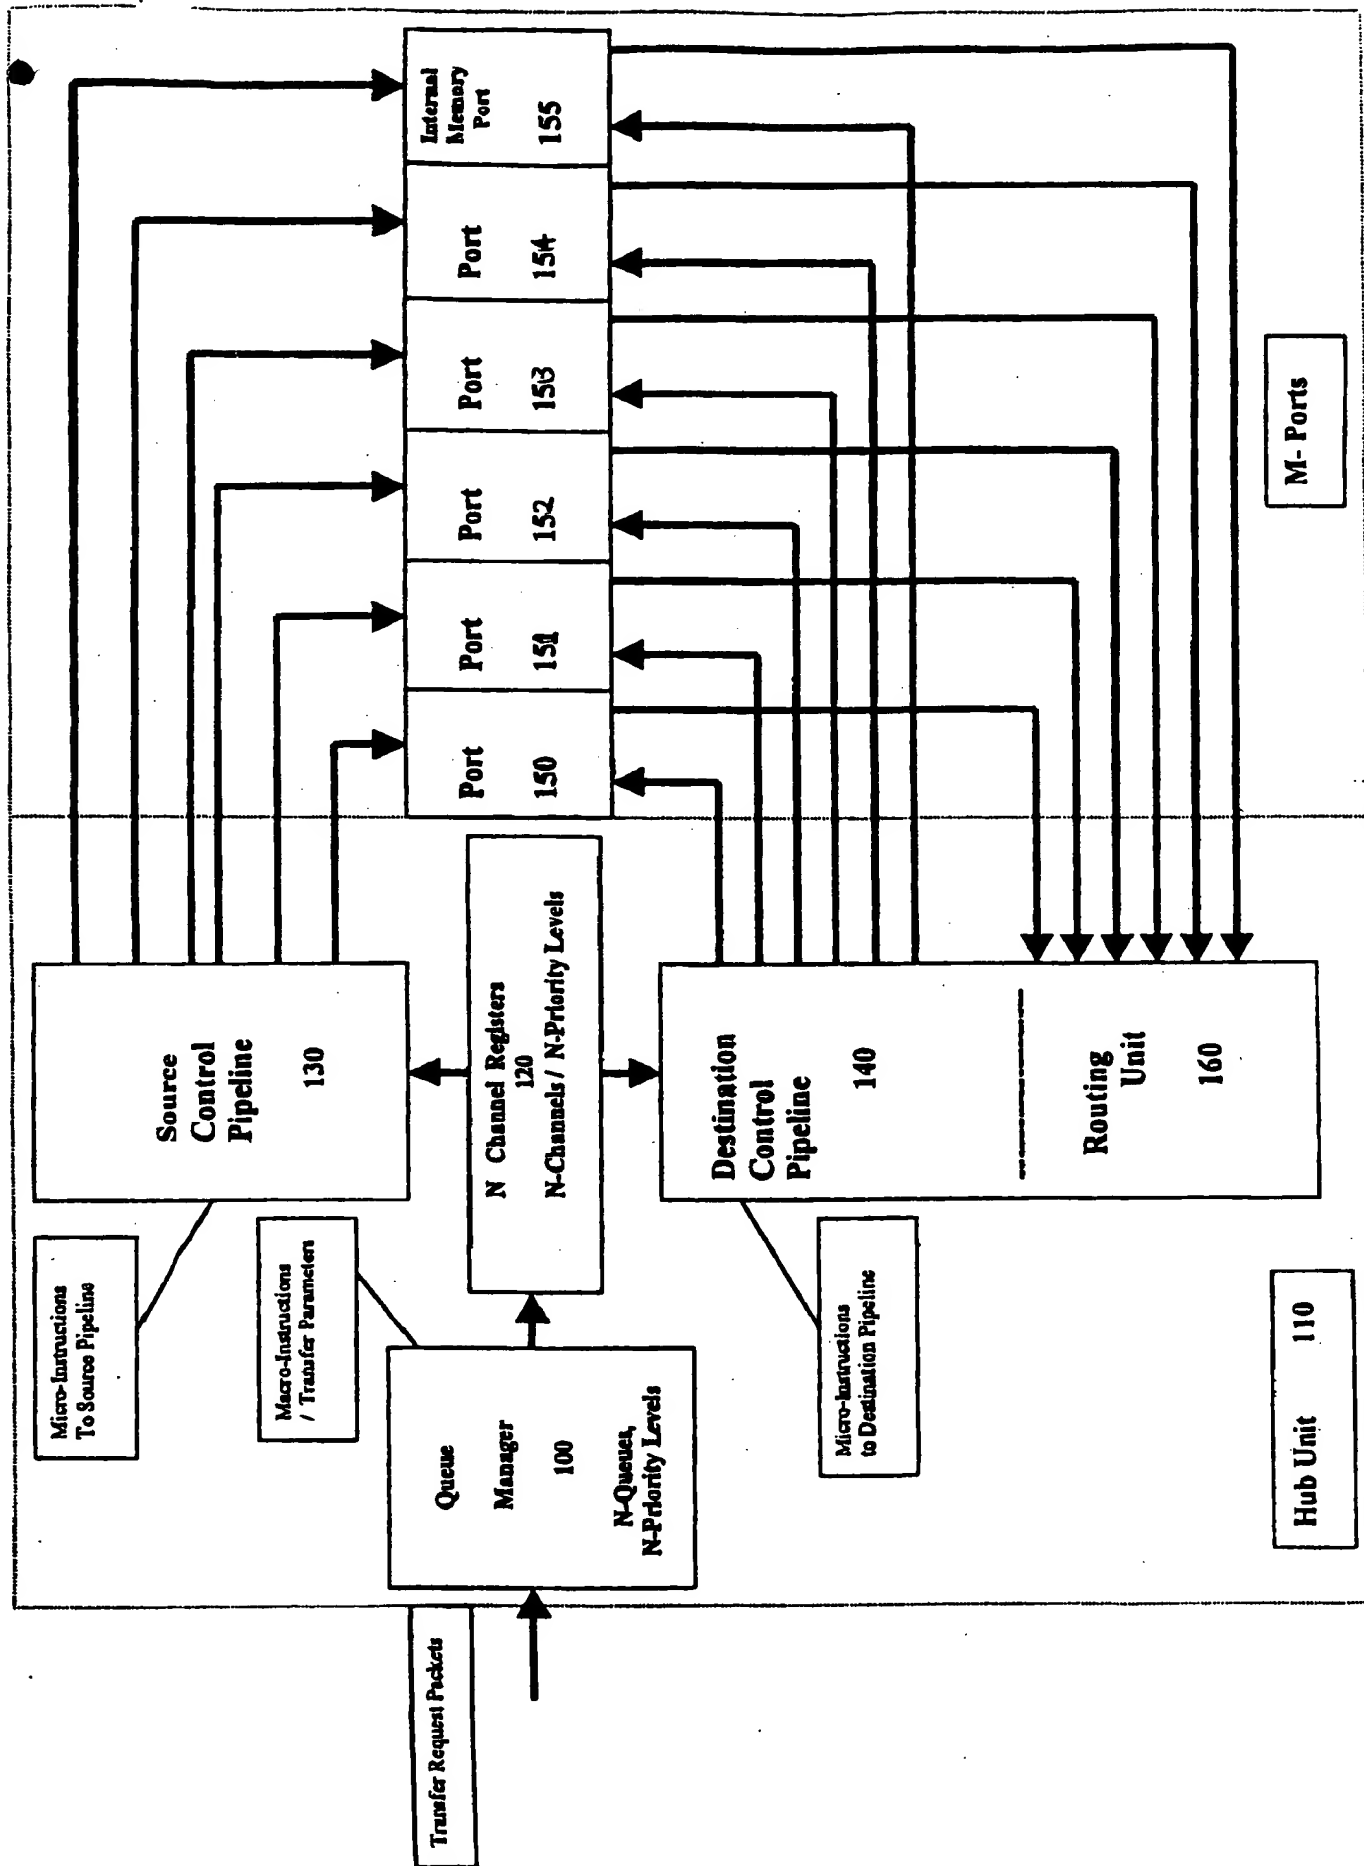
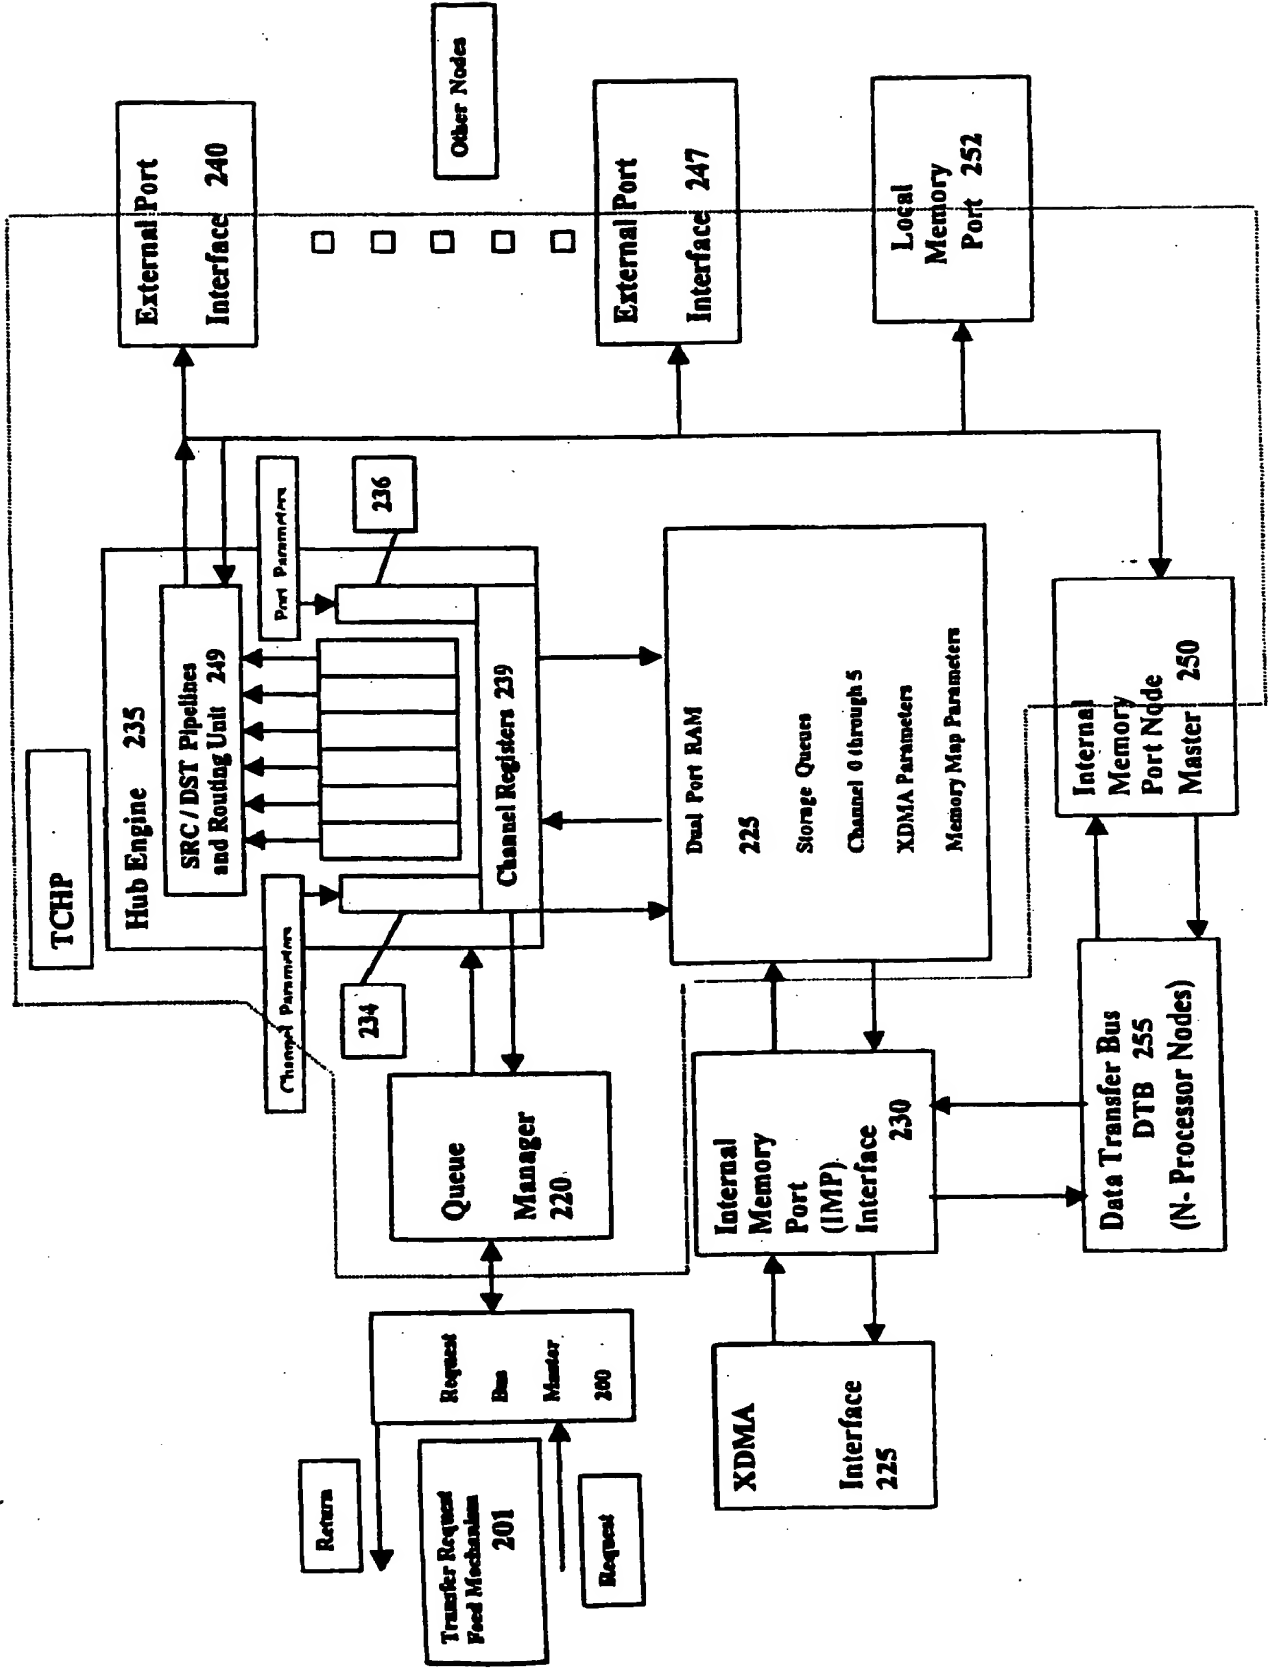


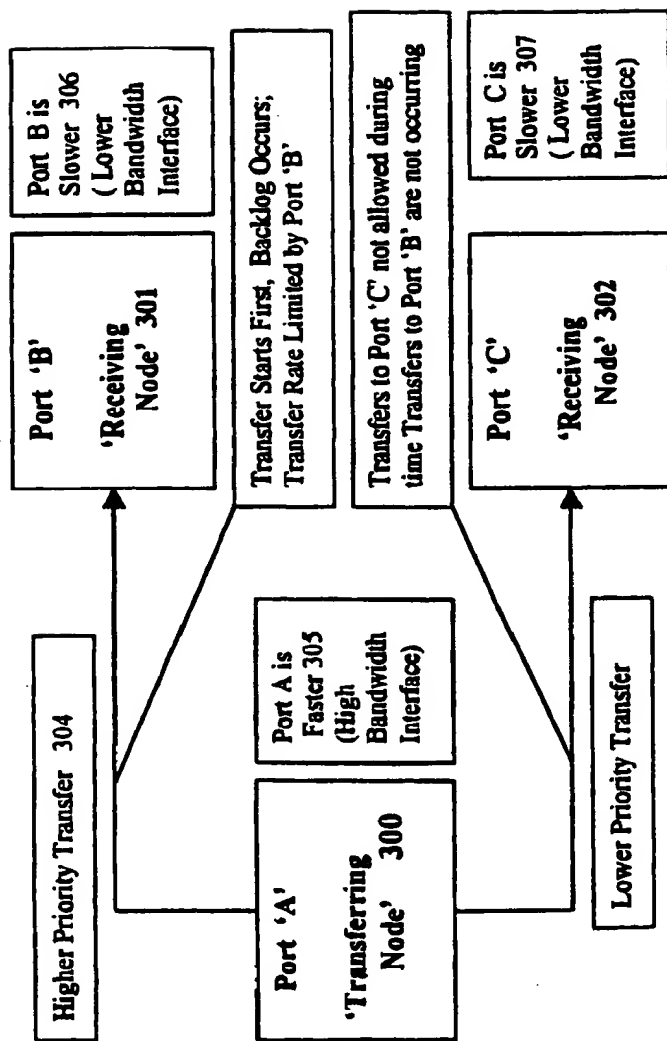
Figure 1

This Page Blank (uspto)

Figure 2



This Page Blank (uspto)



PRIOR
ART

Figure 3

This Page Blank (uspto)

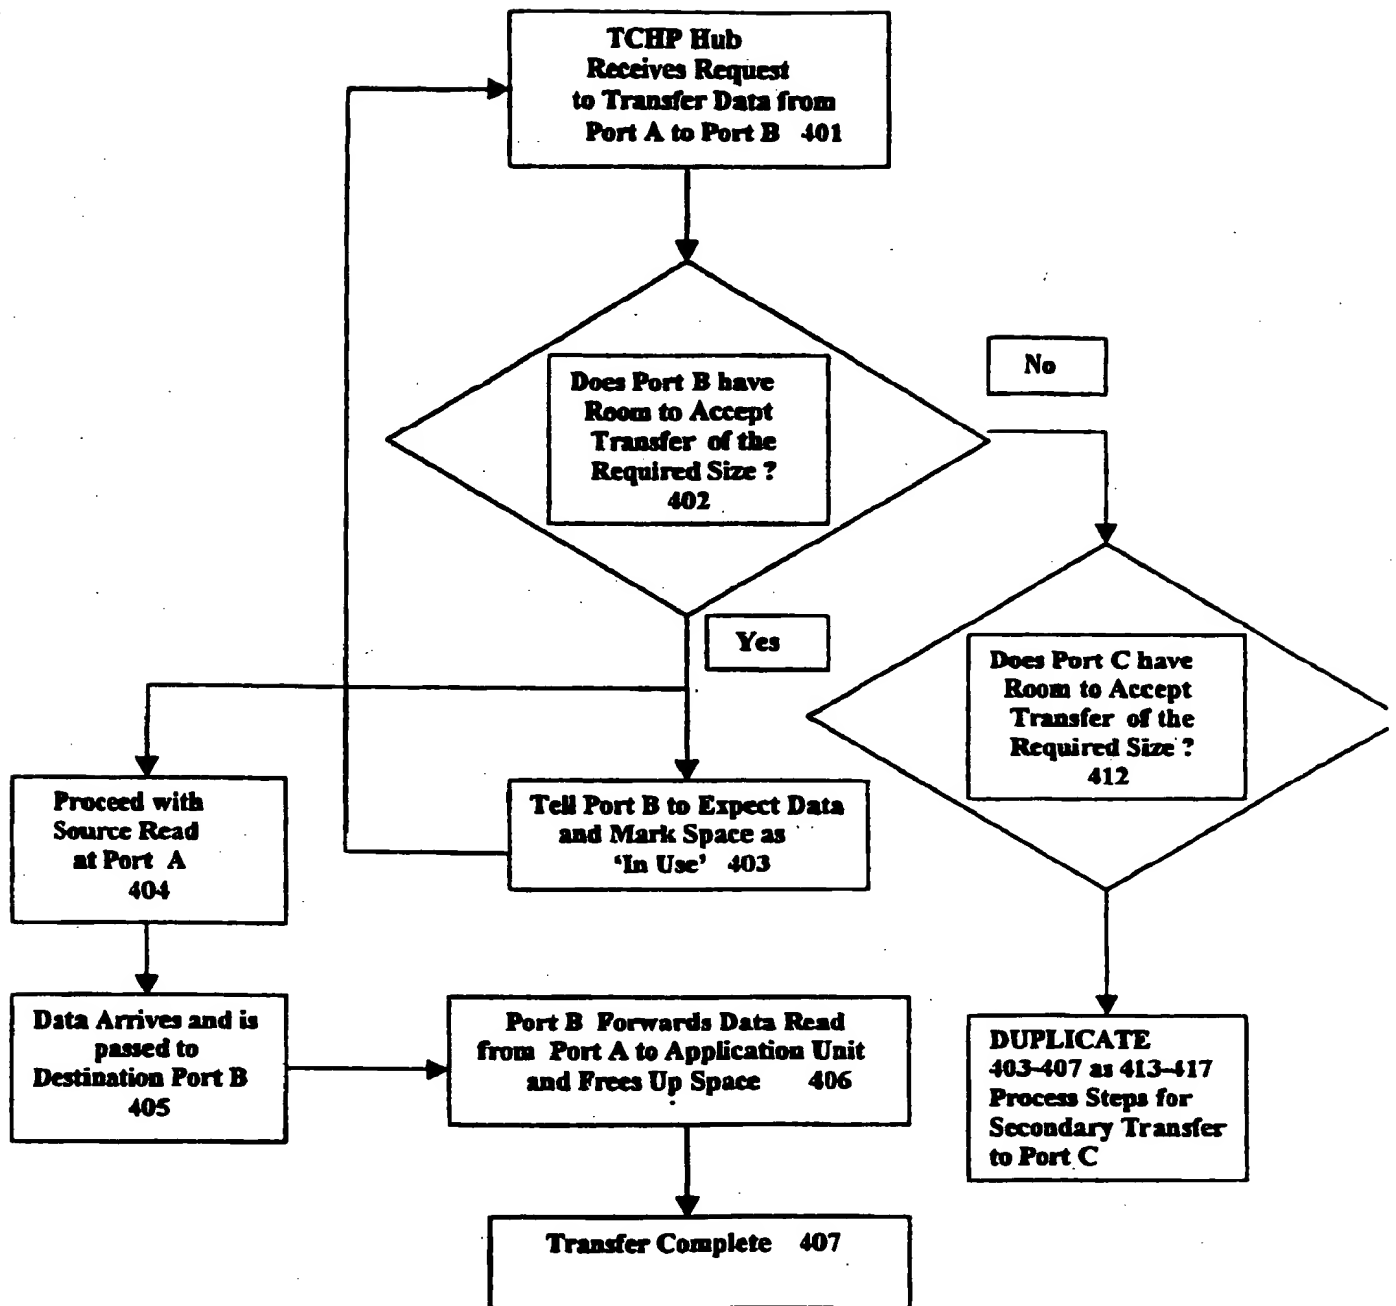


Figure 4

This Page Blank (uspto)

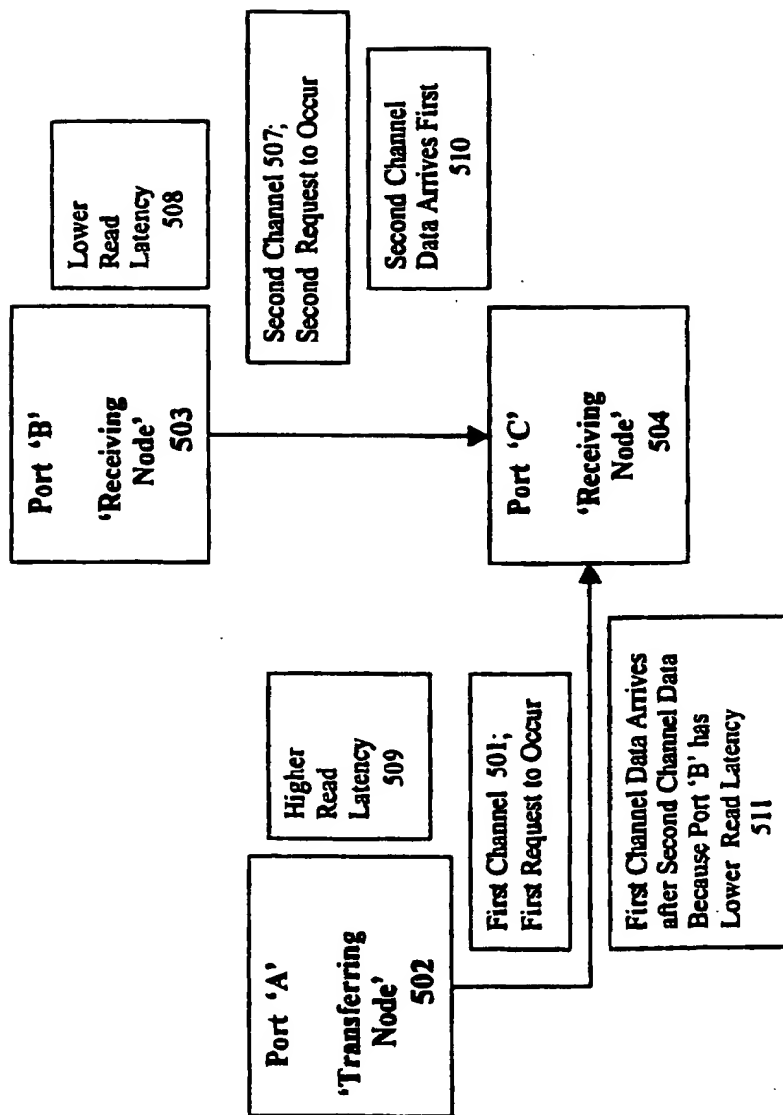


Figure 5

This Page Blank (uspto)

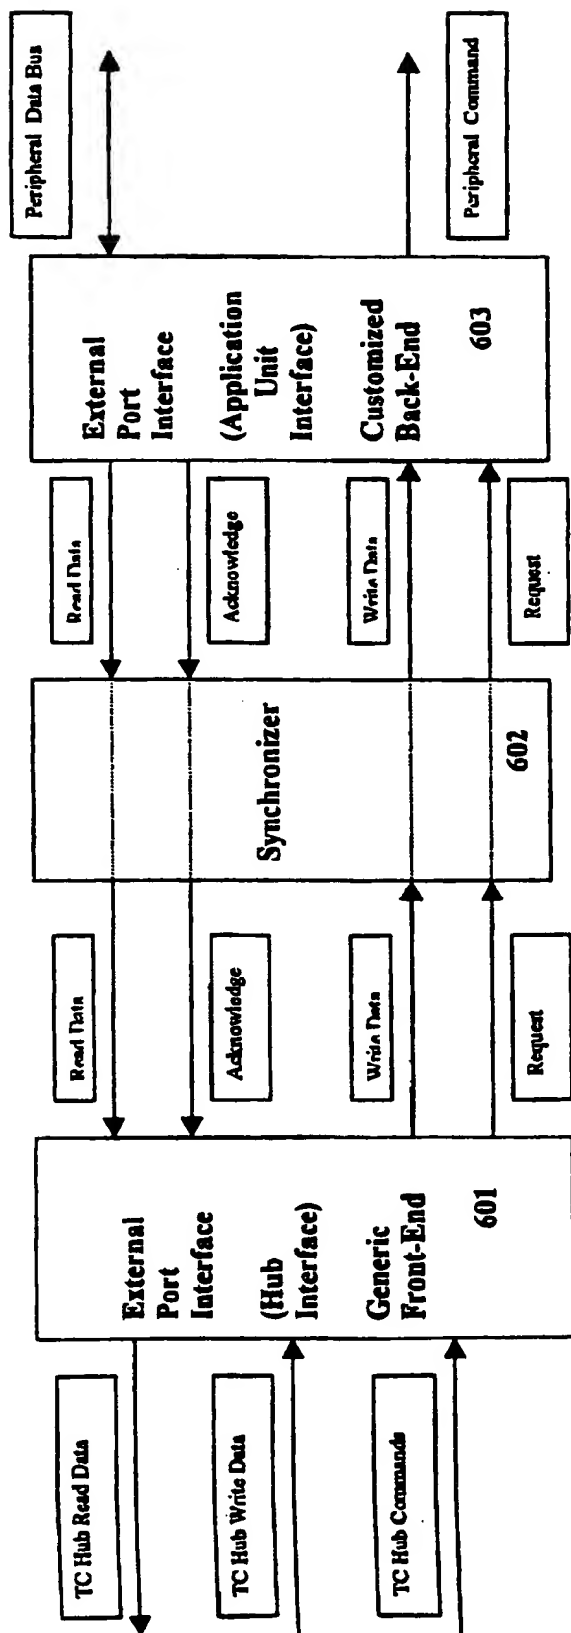


Figure 6

This Page Blank (uspto)

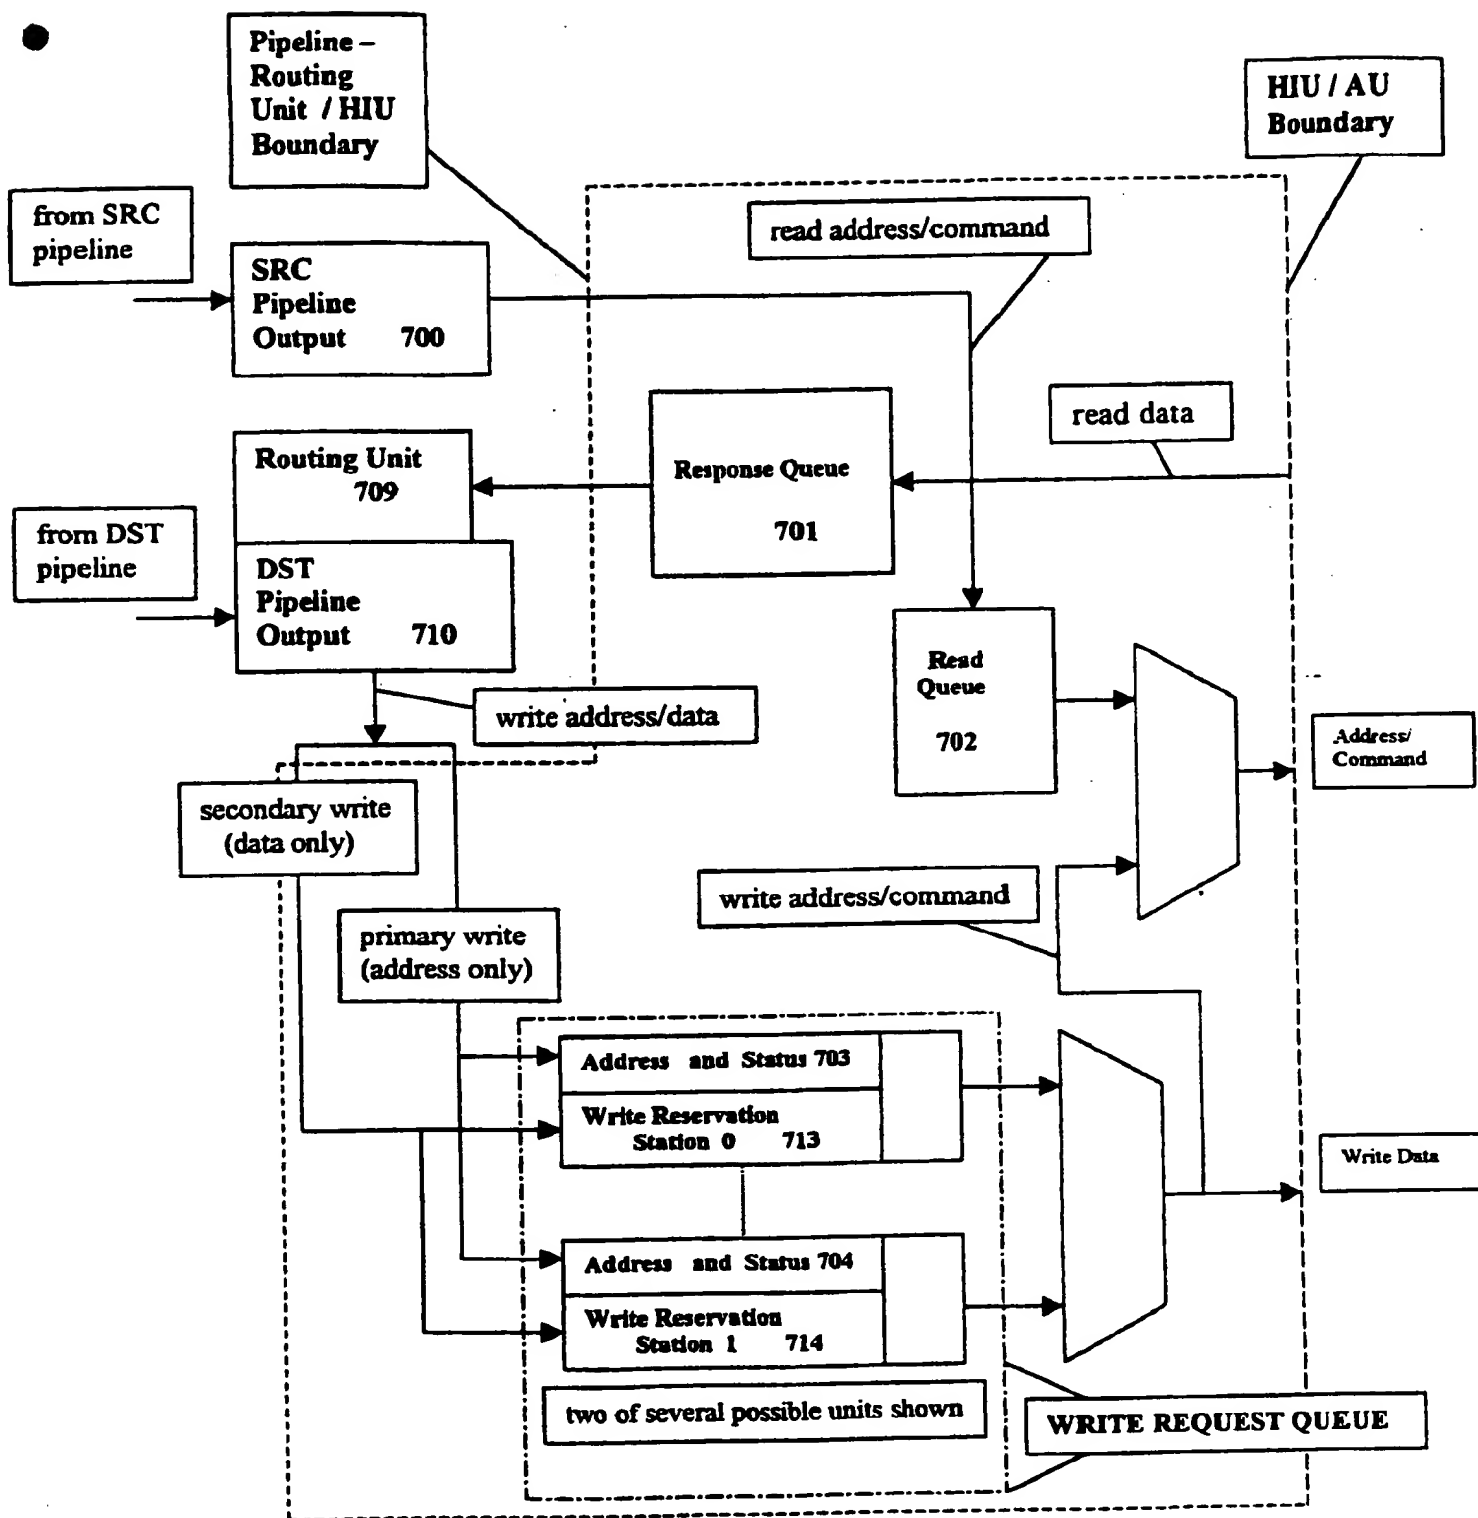
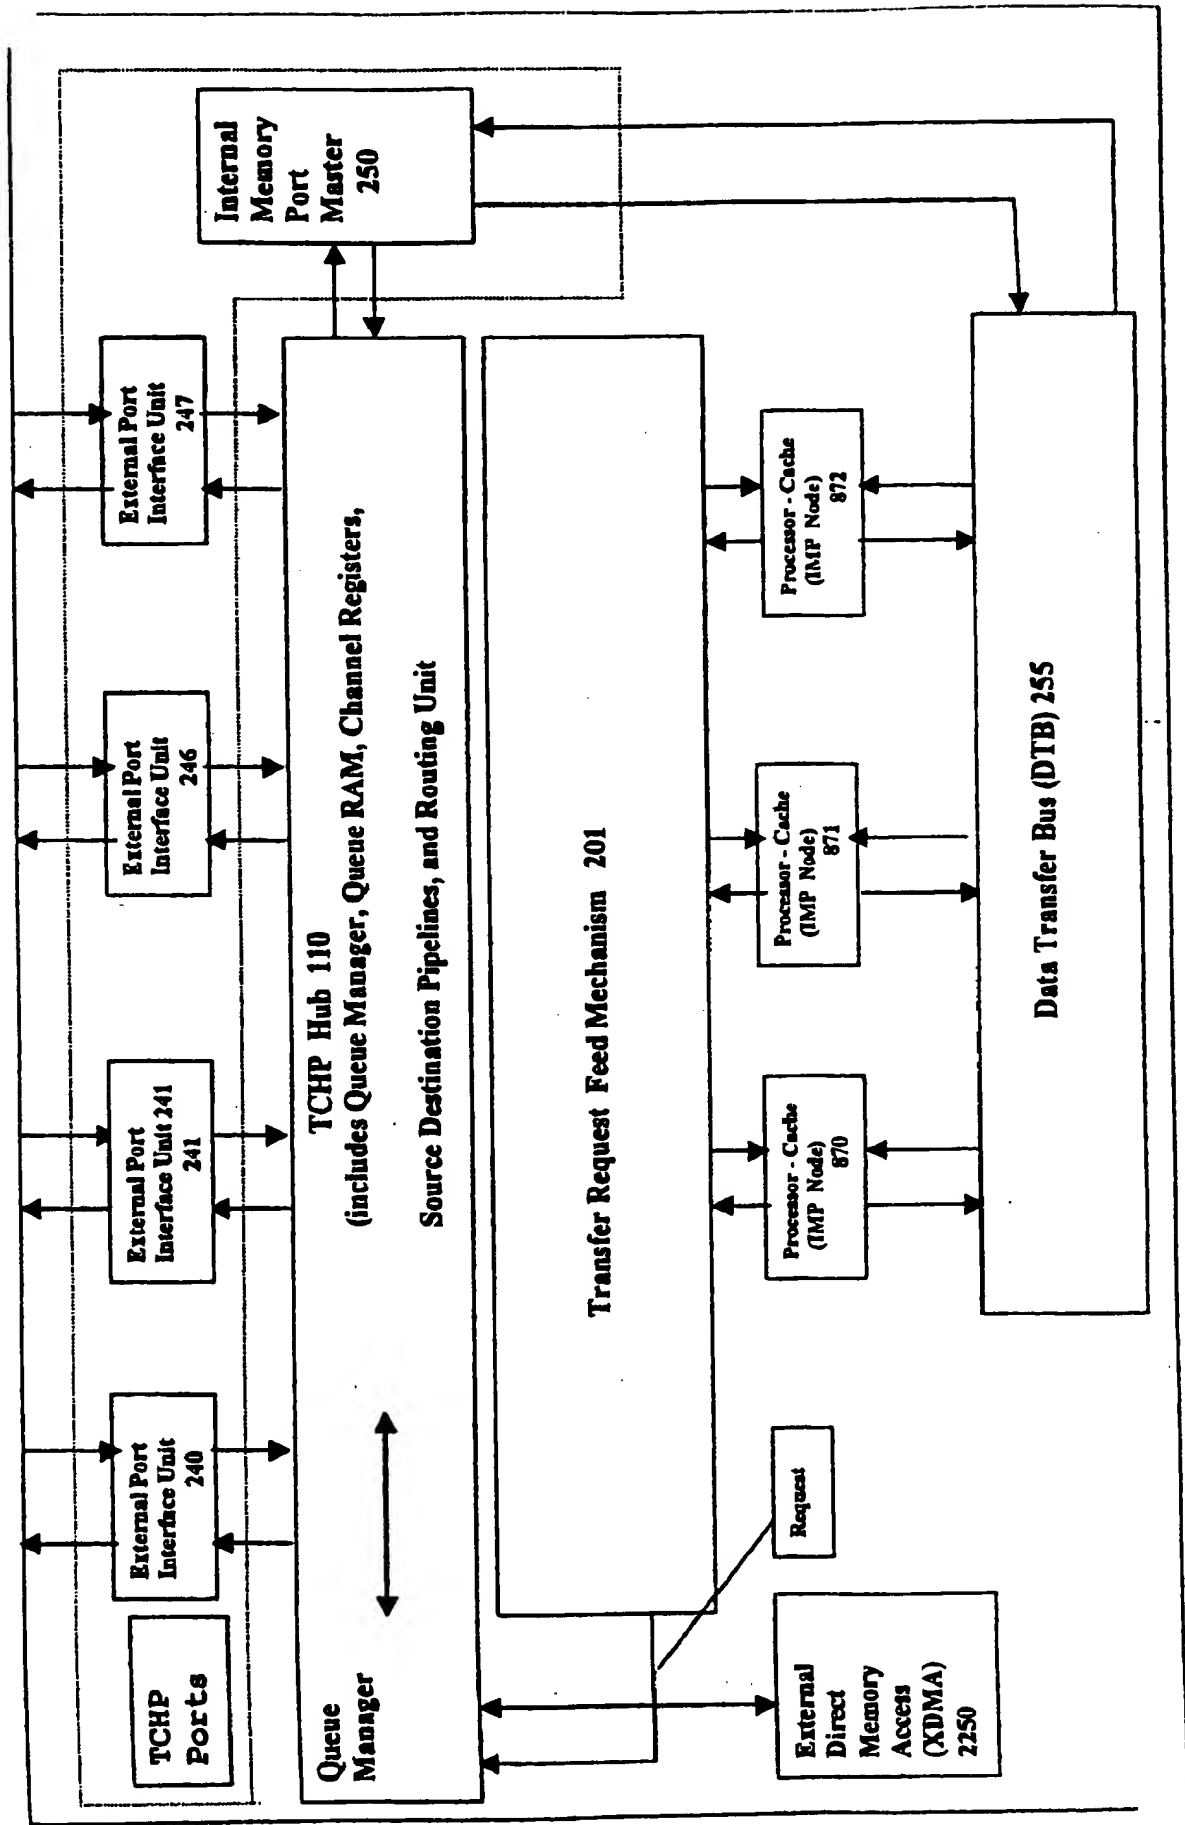


Figure 7

This Page Blank (uspto)

Figure 8



This Page Blank (uspto)

FIG. 9


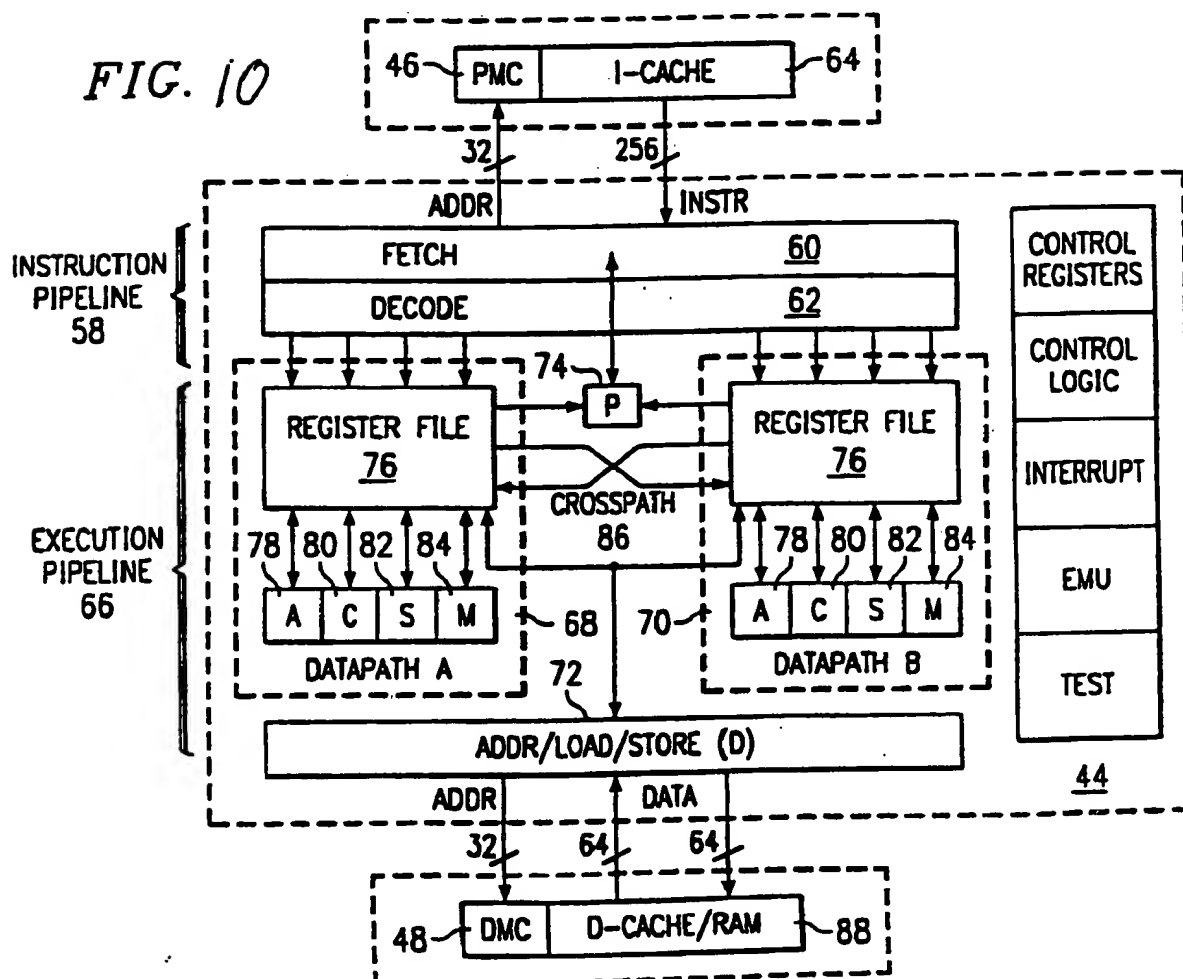


FIG. 10



This Page Blank (uspto)